

Monte Carlo Bayesian inference on a statistical model of sub-gridcolumn moisture variability using high-resolution cloud observations. Part 2: Sensitivity tests and results

Peter M. Norris^{a,b,*} and Arlindo M. da Silva^b

^a*Goddard Earth Sciences Technology and Research, University Space Research Association, Columbia, MD, USA*

^b*Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD, USA*

*Correspondence to: P. M. Norris, Global Modeling and Assimilation Office, NASA/GSFC, Code 610.1, Greenbelt, MD 20771, USA. Email: peter.m.norris@nasa.gov

Part 1 of this series presented a Monte Carlo Bayesian method for constraining a complex statistical model of global circulation model (GCM) sub-gridcolumn moisture variability using high-resolution Moderate Resolution Imaging Spectroradiometer (MODIS) cloud data, thereby permitting parameter estimation and cloud data assimilation for large-scale models. This article performs some basic testing of this new approach, verifying that it does indeed reduce mean and standard deviation biases significantly with respect to the assimilated MODIS cloud optical depth, brightness temperature and cloud-top pressure and that it also improves the simulated rotational–Raman scattering cloud optical centroid pressure (OCP) against independent (non-assimilated) retrievals from the Ozone Monitoring Instrument (OMI). Of particular interest, the Monte Carlo method does show skill in the especially difficult case where the background state is clear but cloudy observations exist. In traditional linearized data assimilation methods, a subsaturated background cannot produce clouds via any infinitesimal equilibrium perturbation, but the Monte Carlo approach allows non-gradient-based jumps into regions of non-zero cloud probability. In the example provided, the method is able to restore marine stratocumulus near the Californian coast, where the background state has a clear swath. This article also examines a number of algorithmic and physical sensitivities of the new method and provides guidance for its cost-effective implementation. One obvious difficulty for the method, and other cloud data assimilation methods as well, is the lack of information content in passive-radiometer-retrieved cloud observables on cloud vertical structure, beyond cloud-top pressure and optical thickness, thus necessitating strong dependence on the background vertical moisture structure. It is found that a simple flow-dependent correlation modification from Riishøjgaard provides some help in this respect, by better honouring inversion structures in the background state.

Key Words: cloud data assimilation; Monte Carlo Bayesian inference; correlation models

Received 18 April 2016; Accepted 19 May 2016; Published online in Wiley Online Library 25 July 2016

1. Introduction

In the Introduction to Part 1 of this series (Norris and da Silva, 2016), we provided motivation for this study by discussing a number of the difficulties associated with the subject of cloud data assimilation (CDA). We particularly emphasized problems associated with the mismatch between the frequently small scales of cloud variability and typical global circulation model (GCM) gridcolumn footprints and the strong nonlinearities present in cloud processes. One of the key problems is that a subsaturated background state cannot produce clouds via any small equilibrium perturbation to moisture. We then provided a detailed description of a new Monte Carlo Bayesian CDA approach designed to

address these problems, with the goal of improving poor model background states to the point where more traditional CDA approaches are able to perform more favourably with them.

This article, Part 2, now discusses the application of the new method and its performance in a number of case studies and sensitivity tests. Section 2 provides an overview, in non-technical terms, of the CDA algorithm described in detail in Part 1. Section 3 sets up a control cloud data assimilation experiment to serve as a baseline to the sensitivity experiments that follow. It is verified that the analyzed cloud state in the control experiment does indeed reflect the Moderate Resolution Imaging Spectroradiometer (MODIS) data that went in: the analyzed cloud optical depth, brightness temperature and cloud-top pressure exhibit

significantly reduced biases with respect to the corresponding MODIS observations. Section 4 examines a number of algorithmic and physical sensitivities of the new method and provides guidance for its cost-effective implementation. This section also demonstrates the significant utility of a simple flow-dependent correlation function modification due to Riishojgaard (1998). Section 5 presents a case study involving marine stratocumulus off the Californian coast. This case is an especially difficult test of the new method, because the observations show stratocumulus very near and sometimes right up to the coast, whilst the background state has a large clear swath in this region. The results are very encouraging: the new algorithm is able to restore the missing stratocumulus, even though the background, subsaturated (clear) regions cannot produce cloudiness for any small equilibrium perturbation to the moisture field. The new algorithm is able to make larger (and non-linearization-based) jumps in parameter space to find an equilibrium cloudy state faithful to the observations. Section 6 presents a short application of the method to assimilation of geostationary cloud retrievals made for the Spinning Enhanced Visible and Infrared Imager (SEVIRI) instrument aboard the *Meteosat-9* platform. The method is easily adaptable to assimilation of these data. A comparison against the assimilation of contemporaneous polar-orbiting *Aqua* MODIS data shows comparable results for both types of observation. Section 7 presents an independent validation of the new method using a non-assimilated dataset, the so-called cloud optical centroid pressure (OCP) retrieved from a rotational–Raman scattering algorithm applied to UV spectral measurements made by the Ozone Monitoring Instrument (OMI) aboard the *Aura* satellite. Finally, Section 8 contains some concluding remarks.

2. Algorithm overview

2.1. Background

Norris and da Silva (2007) conducted a pilot study to investigate the assimilation of International Satellite Cloud Climatology Project (ISCCP) (and later MODIS) and Special Sensor Microwave Imager (SSM/I) cloud data using a parameter estimation technique. Independent Clouds and the Earth's Radiant Energy System (CERES) top-of-the-atmosphere (TOA) cloud forcing estimates were used as a benchmark. The previous generation Goddard Earth Observing System, Version 4 (GEOS-4) model physics used in that study included rather simple diagnostic cloud parametrizations, such as the critical-relative humidity (RH) cloud fraction parametrization of Slingo (1987), and provided an ideal test-bed for assessing the validity of this approach. While GEOS-4 had been carefully tuned to match TOA cloud forcing from CERES, it had excessive cloud amounts in the Tropics and not enough clouds in the extratropics. Not surprisingly, by tuning specific parameters in the Slingo-type scheme, we were able to produce cloud patterns very consistent with the data we assimilated. Having corrected the cloud patterns, the original model tuning now produced much degraded cloud forcing in the long- and short-wave bands. It was not until we constrained the full cloud state with SSM/I cloud water and ISCCP optical-depth estimates that a good agreement with the independent CERES observations could be achieved. The final result was a marked improvement in GEOS-4 long-wave cloud forcing produced by assimilation of the ISCCP cloud mask and SSM/I liquid water path; with additional assimilation of ISCCP cloud optical depth, the short-wave cloud forcing was at least not degraded (Norris and da Silva, 2007). The main lesson learned from this study was that simultaneous observations of cloud fraction, cloud optical depth and cloud condensate were necessary to produce an improved cloud forcing.

While the simple physics in GEOS-4 was a good starting point, we soon realized that, without representing subgrid-scale variability explicitly, we could not take advantage of the wealth of statistical cloud information provided by the high spatial resolution MODIS and other *A-Train* instruments. Fortunately,

an undercurrent of work (e.g. Sommeria and Deardorff, 1977; Smith, 1990; Ricard and Royer, 1993; Xu and Randall, 1996; Larson *et al.*, 2001; Tompkins, 2002; Larson, 2004) has sought a statistical description of subgrid-scale variability of moisture. These statistical schemes specify the subgrid-scale probability density functions (PDFs) of moisture from the outset and the CDA problem becomes one of estimating the parameters of these PDFs.

2.2. Gridcolumn statistical model

While a simple uniform or Gaussian PDF would appear to be a good start, these PDFs fail to model the convectively driven skewed distributions found in nature. In order to model cloud processes, which are strongly weighted to the upper tail of the moisture distribution, more accurately, we adopt here a three-parameter skewed triangular PDF for a single gridbox. This PDF is written in terms of total saturation ratio, $S \equiv q_t/q_s(T)$, defined as the ratio of the *in situ* total moisture content q_t (vapour plus cloud condensate) to the saturation vapour content q_s at the *in situ* temperature T . This PDF is non-zero on (S_L, S_H) , rising linearly from zero at S_L to a mode at S_* and then falling linearly to zero again at S_H (see Figure 1 and also figure 1 of Part 1). A full description of the mathematical properties of this PDF can be found in Appendix A of Part 1. In particular, the gridbox mean water vapour \bar{q}_v , cloud condensate \bar{q}_c and cloud fraction f can be derived from this PDF. Conversely, given $(\bar{q}_v, \bar{q}_c, f)$ arising from this distribution, one can derive the parameters (S_L, S_*, S_H) defining the PDF uniquely. More generally, the macro parameters $(\bar{q}_v, \bar{q}_c, f)$ must satisfy a specific relationship in order for a valid triangular PDF to be obtained. Notice that a two-parameter PDF would pose much more severe constraints on the values of $(\bar{q}_v, \bar{q}_c, f)$.

Similarly to the problem of specifying cloud overlapping for radiative calculations, one also needs a prescription for coupling a vertical stack of triangular PDFs. We adopt here the Gaussian copula formalism of Norris *et al.* (2008). The copula of a set of variables is the joint cumulative distribution function (CDF)




Figure 1. A schematic of a gridcolumn with four layers. Each layer has a skewed triangle PDF in the total saturation ratio $S \equiv q_t/q_s(T)$ (see text). The vertical white line is $S = 1$, representing saturation, with locations to the left being increasingly dry and locations to the right increasingly moist. The dark grey regions of the triangles represent supersaturated air, or cloudiness. The uppermost layer is all subsaturated, or clear. Mean, variance and skewness vary with the layer. The height of each PDF is scaled to the layer height in this schematic (the real PDF height varies such that the PDF area is always one). Each triangle PDF is specified uniquely by a triplet (S_L, S_*, S_H) , representing the lowest, modal and highest S values, respectively.

of the ranks of the variables within their respective univariate margins. The copula can be thought of as a means of producing a multivariate distribution with specified margins. A critical ingredient in this formulation is the specification of a correlation matrix governing inter-layer correlations.

Given this gridcolumn statistical model (GCSM, i.e. layer PDFs plus copula), one can efficiently sample profiles of water vapour and cloud condensate consistent with the model. This is a type of independent column approximation (ICA) subcolumn generation, in the sense that our GCSM does not contain any specification of horizontal correlation, just horizontal PDFs of S coupled only in the vertical.

2.3. Bayesian parameter estimation

Given a set of measurements y (say, 1 km cloud optical depth and cloud-top pressure derived from MODIS), one would like to estimate the GCSM parameters α . Notice that α should include a set of parameters specifying each layer PDF; α may also include the vertical coupling between the layers, for example via one or more vertical decorrelation length-scales associated with inter-layer correlation. Assuming some prior knowledge of α given by the prior PDF $p(\alpha)$, we wish to determine how this knowledge is modified by a vector of observations of the gridcolumn, y . Specifically, we wish to estimate the posteriori PDF $p(\alpha|y)$. This is a straightforward application of Bayes' Theorem:

$$p(\alpha|y) \propto p(y|\alpha)p(\alpha). \quad (1)$$

The first term on the right, $p(y|\alpha)$, is called the likelihood and evaluates the probability of observing y given a parameter state α . Our goal is to estimate $p(\alpha|y)$ to quantify its mode(s), which specifies the most probable α given the observations, as well as some measure of its spread, which indicates the magnitude of the error associated with the modal α estimate. For reasons discussed in Part 1, we use a type of Markov chain Monte Carlo (MCMC) method to characterize $p(\alpha|y)$. This method makes quasi-random jumps around parameter space, such that, as the number of jumps becomes large, the collection of sampled α is consistent with (samples drawn from) the target PDF $p(\alpha|y)$.

2.4. The GEOS-5 Earth System Model

The prior or 'background' state for each of the gridcolumn statistical models is provided by forecasts of the National Aeronautics and Space Administration (NASA) Global Modeling and Assimilation Office (GMAO) GEOS-5 GCM, namely by initialization from the GCM's gridbox mean vapour and condensate contents, \bar{q}_v and \bar{q}_c , and cloud fraction f . More precisely, each three-hourly analysis is initialized from the 'corrector segment' of the incremental analysis update (IAU) of the GEOS-5 data assimilation system. This corrector segment is a 6 h forecast, which gradually introduces the influence of observations from the regular (non-CDA) six-hourly GEOS-5 analysis. Please see section 2a of Rienecker *et al.* (2011) for a more complete explanation.

The model version used in this study is GEOS-5.7.2 and the model's parametrizations are described in Rienecker *et al.* (2011). This version of the GEOS-5 model uses the relaxed Arakawa–Schubert (RAS) cumulus parametrization (Moorthi and Suarez, 1992) for both deep and shallow convection and a prognostic scheme for cloud liquid water and ice, with a PDF-based scheme for large-scale condensation and cloud cover (Bacmeister *et al.*, 2006). As resolution is increased, the convection parametrization is restrained using a stochastic, resolution-dependent limit on deep convection (Tokioaka *et al.*, 1988). At resolutions finer than about 10 km, this scheme essentially prevents deep convection and restricts RAS to act as a shallow convection scheme. For these experiments, we use the current operational resolution, which has a nominal $1/4^\circ$

latitude–longitude grid (more precisely 0.25° latitude by 0.3125° longitude), with 72 vertical layers and a model top at 1 Pa.

Finally, note that each of the eight three-hourly cloud analyses is independent and uses the background state provided by an existing GEOS-5 data assimilation run with an IAU cycle (but without CDA). This study does not include any 'cycling', i.e. the output of the cloud analysis is not used to modify the subsequent GCM forecast. Such a cycling system is certainly planned, but is beyond the scope of the current article.

3. Control run

A number of sensitivity tests have been performed (see below) by perturbing key algorithmic and physical parameters of the Bayesian cloud assimilation system. The control run for these tests is largely defined by Part 1 and specifically with the following choice of parameters.

- (1) $N_{\text{sim}} = N_{\bullet\text{max}} = 64$. Here N_{sim} is the number of simulated subcolumns (per gridcolumn) and $N_{\bullet\text{max}}$ is explained in Appendix B.2 of part 1.
- (2) The number of trials M per point in the multiple-try Metropolis (MTM) chain is a factor $f_M = 1/2$ of the reference value M_* specified in section 2.7.2 of Part 1, i.e. M is about half the number of effectively independent dimensions in the parameter space.
- (3) For our MTM sampler, the optimal Metropolis–Hastings (MH) proposal covariance matrix Σ_q of (13) in Part 1 is amplified by a factor $C = 32$, since the advantage of MTM over MH is that it allows for larger proposal steps.
- (4) An MTM chain length of $n = 200$.
- (5) A vertical correlation length-scale of $L = 100$ hPa.
- (6) The observation-error parametrization for MODIS cloud retrievals is still under development and for this initial phase of our study we consider the perfect observation limit in which observation errors are not added to the simulated observations.

Despite the considerable success of an assimilation run with the above specifications, an examination of the results revealed excessive cloud water path, $CWP = \int \rho_c dz$, in the night-time regions. In retrospect this is not surprising – at night there are no cloud optical thickness (COT) observations to constrain the vertically integrated condensate in each gridcolumn. The remaining observables, T_b and p_c , both saturate very quickly with cloud water path and so cannot constrain CWP . To resolve this issue, for night-time gridcolumns only, we multiply the prior term $p(\alpha)$ of Part 1 by an empirical Gaussian-like term,

$$e^{-C_\lambda [\ln(\lambda/\lambda_B)]^2}, \quad (2)$$

that constrains the ratio of the condensed water path to the total water path, $\lambda = CWP/TWP$, to its triangularized background state value λ_B , where $TWP = \int (\rho_v + \rho_c) dz$. We have found experimentally that the constant $C_\lambda = 256$ works well.

With this modification, the resulting run, denoted b7K64fhC32, forms the control for the sensitivity tests that follow. The run was made for the day of 1 July 2011, comprising eight 3 h cloud analyses (at 0000 UTC, 0300 UTC, ...). Figure 2 shows global plots of COT τ , cloud brightness temperature T_b and cloud-top pressure p_c for the control run for each of the triangularized background (B), observations (O) and result of the Bayesian analysis (A). Note that, in general, the analysis is closer to the observations than the background. We will be more quantitative shortly.

Note that, in all plots and analysis presented in this article, the following applies.

- (1) The use of the term 'background' refers to the triangularized model background state, namely the gridcolumn of skewed triangle PDFs produced by the initialization procedure presented in Part 1.

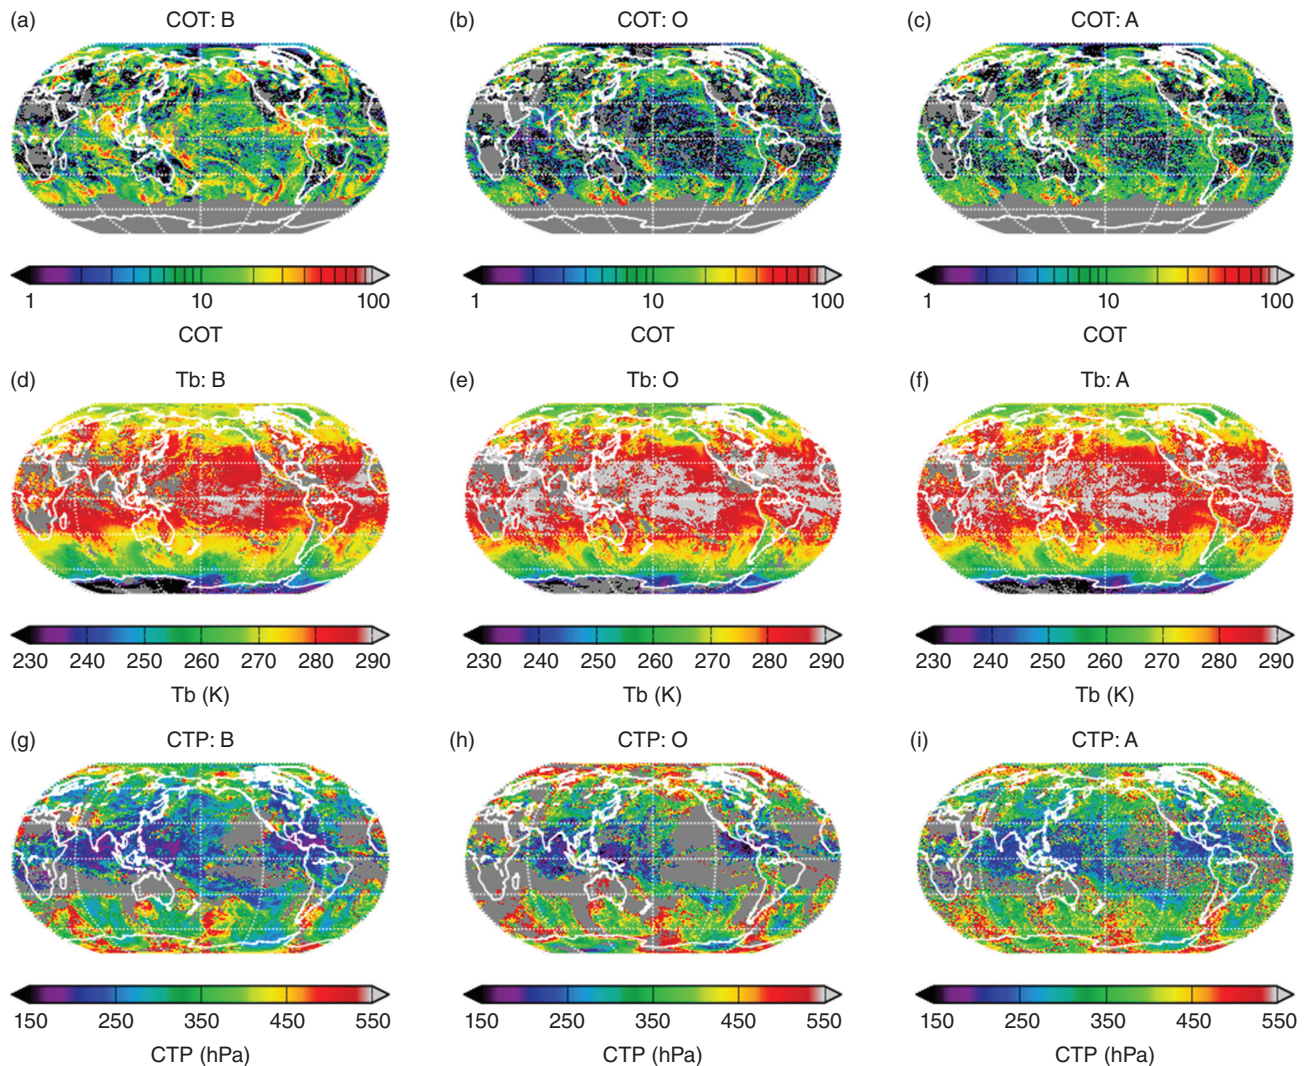


Figure 2. Global plots of (a–c) all-sky gridcolumn mean COT τ , (d–f) in-cloud gridcolumn mean cloud brightness temperature T_b , and (g–i) in-cloud gridcolumn mean cloud-top pressure p_c for the control run b7K64fhC32 for each of the triangularized background (B), observations (O) and analysis (A). Results are averages of the eight three-hourly analyses for 1 July 2011 (at 0000 UTC, 0300 UTC, ...).

- (2) τ is an all-sky gridcolumn mean COT (zeros included for clear pixels/subcolumns) and T_b and p_c are in-cloud gridcolumn means of cloud brightness temperature and cloud-top pressure, respectively.
- (3) T_b is the mean for only the pixels or subcolumns for which T_b is actually used, namely when $p_c > 550$ hPa, and, conversely, p_c is the mean for only $p_c \leq 550$ hPa. τ is the mean over only daytime gridcolumns.

The ‘gridcolumn mean’ here has a specific meaning. For the observations, it is just the mean over the pixels falling within a gridcolumn. For the background and analyzed states, it is the mean calculated over a sample of N_{sim} subcolumns generated at a particular GCSM parameter state α – for the background, that parameter state is the result of the initialization procedure (Part 1, section 2.3.1) and for the analyzed state, that parameter state is the MCMC chain element with the highest posterior probability (Part 1, section 2.7).

Note that each of MODIS *Aqua* and *Terra* provides global coverage every 1–2 days, but the recovery period is much shorter at higher latitudes. When multiple satellites/multiple orbits observe the same gridcolumn within the same time window, the MODIS granule that contributes observations closest to the analysis time is selected. Coverage is also clearly not representative of a diurnal average, given the Sun-synchronous orbits of the satellites. For the average over the eight three-hourly analyses of Figures 2 and 3, the time average is used mainly to allow the global data set to be displayed on a single map, as is common with Level 3 satellite data. However, given the latitude factor as above and the

two satellites, the average is somewhat complicated. Regardless, the observations, background and analysis are averaged in the same way, with a common set of gridcolumns for each time window, to yield a consistent comparison.

Figure 3 shows the corresponding global plots of the O–B and O–A biases. Clearly the analysis biases are much smaller than the background biases. Figure 4 shows the PDFs of O–B and O–A biases in τ for the control run. The peaks of the O–A PDFs are all significantly narrower than for the corresponding O–B PDFs, again indicating the basic success of the cloud assimilation method. Both the modal and the mean biases are also generally smaller in magnitude for O–A than O–B. The figure also shows that the analysis was quite successful in removing unobserved clouds from the background (see figure caption).

Figures 5 and 6 show the corresponding PDFs of biases for the cloud brightness temperature T_b and CO₂-slicing-based cloud-top pressure, p_c . Similar conclusions can be drawn as from Figure 4. Note that the inset panels have a slightly different interpretation than for Figure 4. Namely, for Figure 5 the \bullet refers to the existence of a T_b value, i.e. a ‘low cloud’ ($p_c > 550$ hPa), while the \circ represents not only clear pixels/subcolumns but also those which do not use T_b because the cloud top is too high ($p_c \leq 550$ hPa). Based on this interpretation, the inset illustrates that the fraction of modelled low clouds not seen in the observations is reduced by the analysis (transfer of upper-left red to lower-left blue) and also that the fraction of observed low clouds that are not modelled is also reduced by the analysis (transfer of lower-right red to upper-right blue). Analogous comments apply for Figure 6.

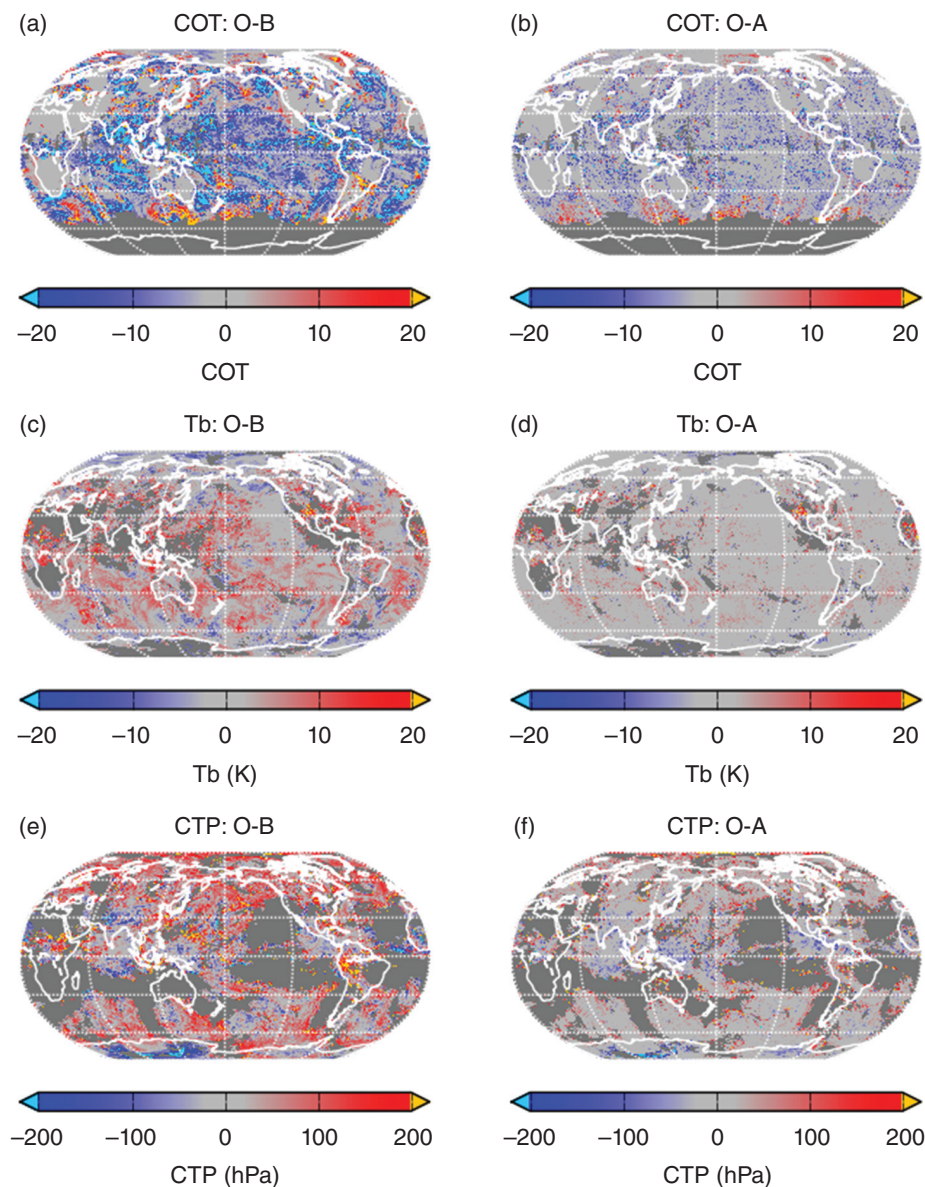


Figure 3. Global plots of O–B and O–A biases for the control run b7K64fhC32 for the observables in Figure 2. Clearly, the analysis biases ((b) COT, (d) cloud brightness temperature, (f) cloud-top pressure) are much smaller than the corresponding background biases (a,c,e).

In view of the clear difference between day and night assimilation (see Eq. (2) and context), we also computed the global bias statistics for the control run separately for day and night gridcolumns. Table 1 shows the results. Note the following: (i) the background biases O–B are significantly smaller in the mean during the day than during the night, but very similar in standard deviation and (ii) the correction of mean bias in T_b and p_c by the cloud analysis is significantly stronger during the night than during the day; the same also applies for the standard deviation in the bias, though to a much lesser degree. It would seem that the daytime analysis, with its additional constraint on the τ observable, achieves a reduced alignment to the T_b and p_c observables and that addition of the night-time prior constraint on CWP/TWP , although having a similar effect, is not as dominant.

4. Sensitivity tests

The above results simply show that assimilating observations moves a background state to an analysis with a better fit to those observations. This is to be expected, and only acts as a gross check of the assimilation system. It does not necessarily imply that the analysis is *better* according to other measures, such as internal dynamic/thermodynamic consistency or vertical structure not constrained by the observations. We will now perform various

sensitivity tests on algorithmic and physical parameters of the CDA system and will judge these sensitivity tests according to how they impact the fit to observations. Again, a better fit does not imply a better analysis. Rather, our confidence in the analysis must be built by examining it in various case studies and by validating it with additional non-assimilated observables. We will begin this additional validation with a marine stratocumulus case study and an OCP validation later in the article.

Table 2 lists the global mean biases in τ , T_b and p_c and the standard deviation (sdev) in the biases for the control run b7K64fhC32 and for several sensitivity experiments varying key algorithmic parameters. The experiment b7K64fhC16 reduced the proposal covariance by half with $C = 16$ and produces a slightly worse mean and sdev O–A than the control. Tests with $C = 64$ (not shown) produced mixed results when compared with the control, but no significant improvement. The experiment b7K64fhC32 uses twice as many trials per MTM chain point as the control, namely $f_M = 1$. We expected this to be better, since using the same number of trial points as there are effectively independent dimensions in the parameter space seems reasonable. Indeed, the results are slightly better, but certainly not enough to justify twice the expense of this experiment (doubling the number of trials approximately doubles the experiment timing). The experiment b7K128fhC32 doubles N_{sim} (and N_{max}) to 128. On the whole, it is slightly




Figure 4. PDFs of the biases O–B (red) and O–A (blue) in all-sky gridcolumn mean cloud optical thickness for control run b7K64fhC32. Biases are shown for six regions (a)–(f), as noted, where the latitude boundaries between Tropics, mid and high latitudes are at 22.5 and 50°. The peaks of the O–A PDFs are all significantly narrower than those of the corresponding O–B PDFs, indicating the basic success of the cloud assimilation method. Both the modal biases (location of the peak) and the mean biases (dashed vertical lines) are also generally smaller in magnitude for O–A than O–B. The mean and standard deviation of the biases are also explicitly printed in the upper left of each panel. Note that all these results are for the subset of gridcolumns that are non-clear (either partially cloudy or overcast) for both the observations (in the gridcolumn) and the model (B or A). This subset is also represented in the ●● quadrant of the upper right inset of each panel. This inset shows the four-way fractional split between clear (○) and non-clear (●) cases for the observations O and the model M (B or A). The fact that there is a transfer of probability from cloudy B to clear A when the observations are clear (i.e. red in upper left to blue in lower left quadrants) is also indicative of the success of the cloud assimilation method.

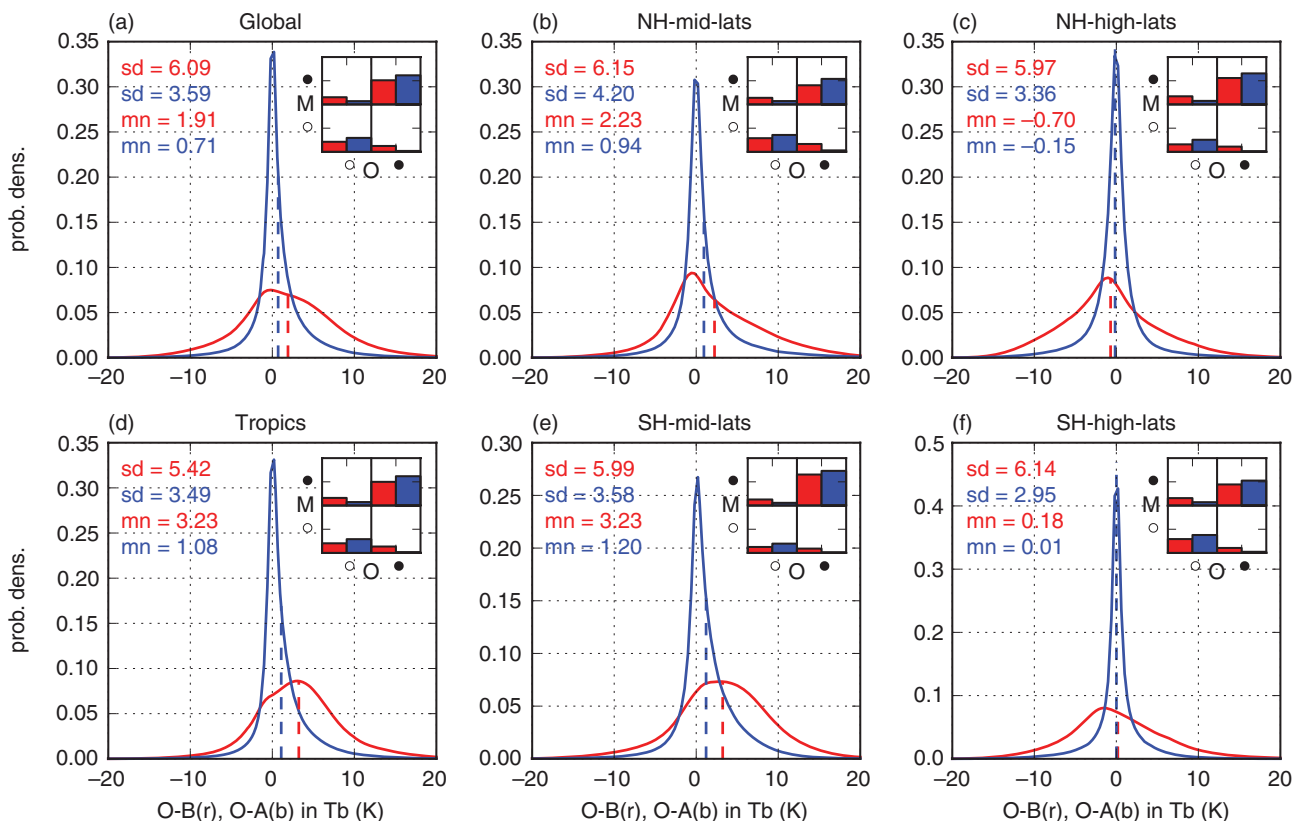


Figure 5. As in Figure 4, but for brightness temperature T_b and including night-time as well as daytime gridcolumns. See the text for important comments on the inset, which has a different interpretation than in Figure 4.

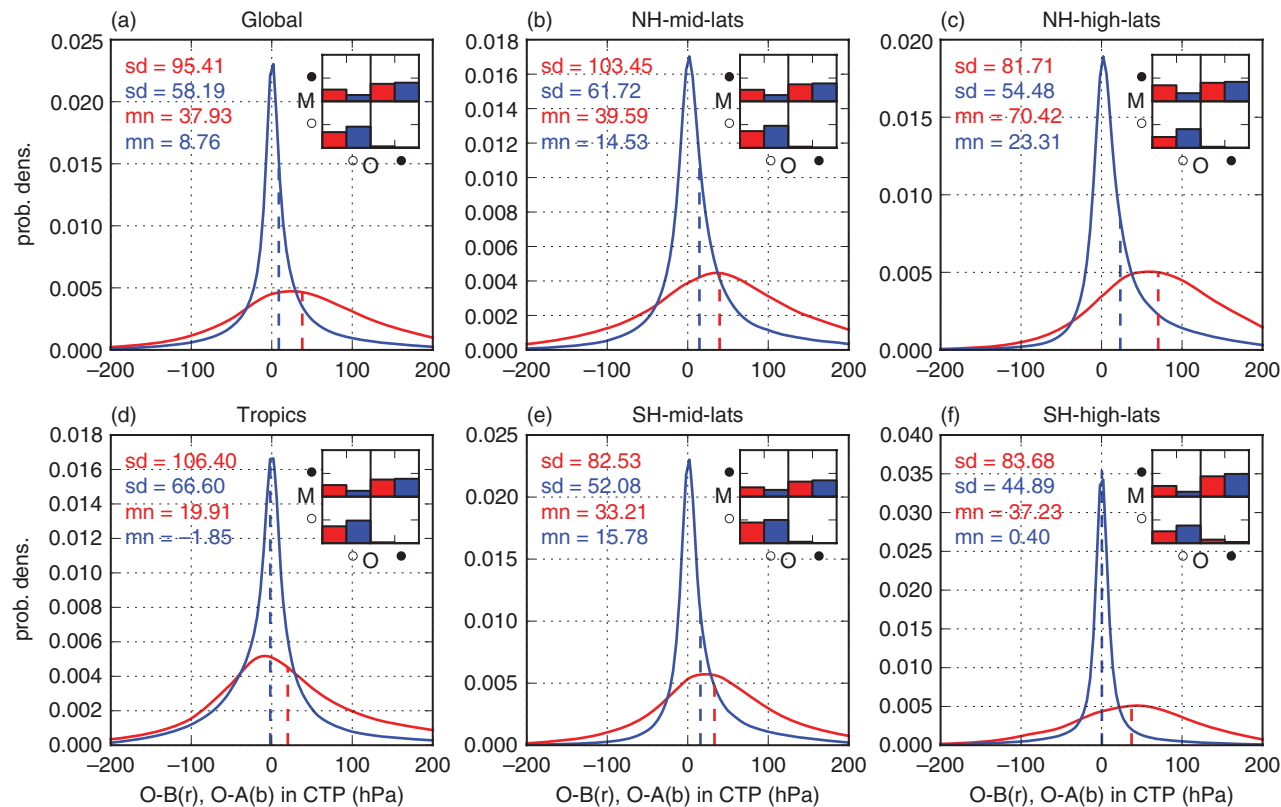


Figure 6. As in Figure 5, but for CO₂-slicing-based cloud-top pressure. See the text for important comments on the inset.

Table 1. Summary of the global biases for the control run b7K64fhC32 and the control run with day and night subsetting. The format of each box is O-B → O-A, where the mean bias is reported in the upper section and the standard deviation of the bias in the lower section. No column for τ bias is reported, since τ is only available during the day.

Mean	T_b	p_c
b7K64fhC32	1.91 → 0.71	37.93 → 8.76
Day only	1.48 → 0.96	32.17 → 14.39
Night only	2.21 → 0.54	41.54 → 5.25
sdev	T_b	p_c
b7K64fhC32	6.09 → 3.59	95.41 → 58.19
Day only	6.04 → 4.13	97.12 → 62.43
Night only	6.12 → 3.16	94.14 → 55.10

better, but again not enough to justify twice the expense. The experiment b7K64fhC32.n400 doubles the length of the MTM chain to 400 elements. It is also slightly better than the control, but yet again not worth twice the run-time expense. All the above experiments involve algorithmic perturbations and we conclude that the control b7K64fhC32 is satisfactory and cost-effective.

One final algorithmic sensitivity experiment, b7G64fhC32, replaced the kernel density estimate (KDE) method of likelihood evaluation, described in Part 1, section 2.6, with a simple 2D Gaussian likelihood evaluation, also discussed in that section. While this Gaussian likelihood forces the likelihood functions $\hat{p}_{p_c}((\ln \tau, p_c)|\alpha)$ and $\hat{p}_{T_b}((\ln \tau, T_b)|\alpha)$ to be unimodal and of Gaussian form (e.g. with elliptic contours), the results from Table 2 are actually a small improvement on the control in all measures except global p_c bias. Although we retain the KDE likelihood in the control for its generality and all-round good performance, we intend to continue to investigate the Gaussian likelihood, because the simplicity of the Gaussian may allow some simplification and speed-up of the code (and perhaps a simpler integration of our CDA method with more traditional 3D- and 4D-Var approaches). That being said, the

Table 2. Summary of the variation in global biases between the control run b7K64fhC32 and different algorithmic sensitivity experiments. The format of each box is O-B → O-A, where the mean bias is reported in the upper section and the standard deviation of the bias in the lower section.

Mean	τ	T_b	p_c
b7K64fhC32	-4.82 → -2.18	1.91 → 0.71	37.93 → 8.76
b7K64fhC16	-4.82 → -2.17	1.90 → 0.79	37.96 → 9.92
b7K64fhC32	-4.83 → -1.87	1.91 → 0.66	37.91 → 6.89
b7K128fhC32	-4.78 → -1.99	1.96 → 0.67	37.49 → 6.11
b7K64fhC32.n400	-4.82 → -1.86	1.91 → 0.64	37.92 → 6.80
b7G64fhC32	-4.82 → -1.94	1.91 → 0.62	37.95 → 10.41
sdev	τ	T_b	p_c
b7K64fhC32	15.50 → 7.02	6.09 → 3.59	95.41 → 58.19
b7K64fhC16	15.51 → 7.23	6.09 → 3.81	95.38 → 61.15
b7K64fhC32	15.51 → 6.64	6.09 → 3.49	95.36 → 56.92
b7K128fhC32	15.40 → 6.44	6.08 → 3.39	95.27 → 55.87
b7K64fhC32.n400	15.50 → 6.47	6.09 → 3.41	95.38 → 54.93
b7G64fhC32	15.50 → 6.66	6.09 → 3.49	95.36 → 54.21

Gaussian likelihood version is not currently faster than the KDE version, because the statistical subcolumn generation dominates the cost of evaluation of the likelihood at the observations. This might be reversed with further coding improvements. (We also plan to try a Gaussian copula (GCOP) likelihood, a middle ground between the KDE and Gaussian likelihoods, as mentioned in Part 1. This is distinct from the GCOP model of layer overlap already included in this study.)

Table 3 examines the sensitivity of global biases to more physical parameters. The experiment b7K64fhC32.L200 doubles the vertical correlation length-scale L to 200 hPa. This is somewhat worse than the control. It is technically possible to include L in the parameter list for Bayesian inference, but we have not yet tried this. The experiment b7K64fhC32.pTrop replaces the fixed $p_{\text{ramp}} = 100$ hPa of Part 1, section 2.3.3, with a variable tropopause pressure as output by the background GEOS-5 simulation and replaces the fixed $p_{\text{lim}} = 50$ hPa with $p_{\text{ramp}} - 50$ hPa. (Note: p_{ramp}

Table 3. Summary of the variation in global biases between the control run b7K64fhC32 and sensitivity experiments examining physical parameters. The format is the same as in Table 2.

Mean	τ	T_b	p_c
b7K64fhC32	−4.82 → −2.18	1.91 → 0.71	37.93 → 8.76
b7K64fhC32.L200	−4.40 → −3.61	1.79 → 1.15	40.95 → 17.40
b7K64fhC32.pTrop	−4.81 → −2.16	1.91 → 0.66	37.90 → 6.26
b7K64fhC32.lnP	−4.64 → −2.21	1.99 → 1.05	37.81 → 6.66
b7K64fhC32.lnPtr	−4.62 → −2.34	1.99 → 1.04	37.84 → 5.31
b7K64fhC32.RiiS1	−5.12 → −0.66	2.09 → 0.35	35.15 → 3.05
sdev	τ	T_b	p_c
b7K64fhC32	15.50 → 7.02	6.09 → 3.59	95.41 → 58.19
b7K64fhC32.L200	15.26 → 9.07	6.17 → 4.29	95.13 → 67.50
b7K64fhC32.pTrop	15.49 → 7.12	6.09 → 3.55	95.37 → 56.94
b7K64fhC32.lnP	15.40 → 7.02	6.14 → 3.80	95.21 → 55.61
b7K64fhC32.lnPtr	15.39 → 7.29	6.14 → 3.82	95.21 → 54.95
b7K64fhC32.RiiS1	15.70 → 6.22	6.05 → 3.23	95.68 → 53.65

is the lowest pressure at which full parameter variability is permitted. Above this level, the parameter variability decays linearly with pressure to zero at p_{lim} . This ramp affects both the \bar{S} prior of Part 1, section 2.3.3, and the MTM estimated target covariance, Σ_π , discussed in Part 1, section 2.7.2). The results for b7K64fhC32.pTrop are very similar to the control run, with the exception that the mean cloud-top pressure bias is significantly reduced (from a 77% reduction in O–B → O–A to an 83% reduction). This improvement makes sense: we are constraining the analyzed cloud-top pressure to approximately below the background tropopause, rather than a nominal level of 100 hPa in the control. In retrospect, this modification should have been included in the control and will be for future work, but not for the current article, except where noted.

The experiment b7K64fhC32.lnP uses $\ln p$ rather than p in matters relating to the specification of vertical correlation. This has some justification, since $\ln p$ is a proxy for the height above the surface, z , and it is more common to parametrize vertical correlation in terms of z than p . The implementation is as follows: in the evaluation of the vertical correlation matrix \mathbf{C} in Part 1, Eq. (3), $\xi_{kk'} = |p_k - p_{k'}|/L$ is replaced by $\xi_{kk'} = |\ln p_k - \ln p_{k'}|/L_{\ln}$, where L_{\ln} is a ‘length-scale’ in $\ln p$, given by $L_{\ln} = \ln[(p_{\text{ref}} + \Delta p/2)/(p_{\text{ref}} - \Delta p/2)]$, such that $p_{\text{ref}} + \Delta p/2$ and $p_{\text{ref}} - \Delta p/2$ are separated by one length-scale in $\ln p$, where p_{ref} is some reference pressure, set to 500 hPa for this experiment, and $\Delta p = 100$ hPa, akin to the $L = 100$ hPa of the control. This means that 450 and 550 hPa have the same decorrelation in both the control b7K64fhC32 and the b7K64fhC32.lnP experiment, but that other levels separated by 100 hPa now have different decorrelations from the control. In particular, upper levels separated by 100 hPa are relatively less correlated, since Δp higher in the atmosphere represents a relatively greater altitude difference than the same Δp near the surface. In addition, the linear ramp in p between p_{ramp} and p_{lim} discussed above for near-tropopause decay of parameter variability now becomes linear in $\ln p$. Lastly, in the calculation of the number of trials M per point in the MTM chain (Part 1, section 2.7.2), $(p_0 - p_{\text{lim}})/L$ is replaced by $\ln(p_0/p_{\text{lim}})/L_{\ln}$. The performance of the b7K64fhC32.lnP experiment is at best mixed, with small improvements for p_c and degradations for T_b . The lnP experiment is also about 60% more expensive than the control, due to the increase in the number of trials M per MTM chain point, as described above $[\ln(p_0/p_{\text{lim}})/L_{\ln} \div (p_0 - p_{\text{lim}})/L = \ln(1000/50)/\ln(550/450) \div 950/100 \approx 15 \div 9.5 \approx 1.6]$. Using a p_{ref} lower in the atmosphere exacerbates this problem. The experiment b7K64fhC32.lnPtr is a combination of b7K64fhC32.pTrop and b7K64fhC32.lnP. The pTrop modification generally lowers p_{lim} within the atmosphere and therefore decreases M , resulting in a small overall decrease in computation time for b7K64fhC32.lnPtr relative to the control. However, the results from Table 3 are not, on balance, better than the lnP

experiment, with a slight improvement in p_c being offset by a slight deterioration in COT.

Finally, the experiment b7K64fhC32.RiiS1 implements a version of the simple flow-dependent correlation function proposed by Riishojgaard (1998). The idea is to replace the ‘pressure distance’ based vertical correlation matrix \mathbf{C} of Part 1, Eq. (3), by the Hadamard product of itself and another correlation matrix \mathbf{D} with elements

$$D_{kk'} = \nu(|\theta(\alpha_k^b) - \theta(\alpha_{k'}^b)|),$$

where ν is a correlation function on \mathbb{R} and θ is a function of the background state vector α^b for a layer. In this study, we use the Gaussian correlation function $\nu(r; \sigma) = \exp[-r^2/(2\sigma^2)]$. Riishojgaard suggests the use of some conservative function for θ , on the basis that Lagrangian transport preserves conservative properties. In effect, the \mathbf{D} matrix would therefore decorrelate regions that have different Lagrangian origins. The use of a conservative θ is most important for horizontal correlations, since advection-based variability dominates in the horizontal. In the vertical, other processes such as turbulent diffusion and radiative heating become significant. For our case of vertical correlations, we could use the conservative total water content, q_t , but we can just as well use the control variable \bar{S} , the mean total saturation ratio, which has a smaller dynamic range but still captures well the sharp humidity inversion at the top of the planetary boundary layer, for example. This is really our goal – to use the \mathbf{D} term to provide strong decorrelation across inversion features in the vertical, which are otherwise not decorrelated adequately by the fixed pressure scale L of the control \mathbf{C} model. Thus, the experiment b7K64fhC32.RiiS1 applies the above \mathbf{D} modification, with $\theta(\alpha_k^b) = \bar{S}_k^b$ and a standard deviation $\sigma = \kappa_S \sigma_{\bar{S}}$, where $\sigma_{\bar{S}} = 0.1$ is the prescribed standard deviation of \bar{S} in the control (see Part 1, section 2.3.3) and κ_S is a constant, set to one for this experiment. Looking at the results from Table 3, we see that this experiment is superior to the control in all measures. We believe that using the \mathbf{D} matrix honours better the basic vertical structure of the atmosphere, as represented in the background state, and, to the extent that this background state is able to capture realistic vertical structure, is producing an analyzed state more consistent with the observations. In retrospect, this modification should have been included in the control and will be for future work, but not for the current article, except where noted.

5. California stratocumulus study

Figure 7 shows a case study off the west coast of North America, in which marine stratocumulus is present right up to the coast in Southern California and Baja California, but is absent in a wide swath off this coast in the background. In such cases, it is typically very difficult for a cloud data assimilation system to restore equilibrium cloud to such regions because the background, subsaturated clear regions cannot produce cloudiness for any small equilibrium perturbation to the moisture field. However, the Monte Carlo Bayesian system is able to restore the stratocumulus, since it makes non-gradient-based jumps in parameter space. The cloud brightness temperature T_b , in the restored region is in good agreement with the observations, although the COT appears a little too high. Note that these results are for the experiment b7K64fhC32.RiiS1, i.e. the control with the Riishojgaard (1998) correlation modification. While the pure control was almost as good and certainly also restored the near-coast stratocumulus, it produced a more noisy analysis with even higher COT in the restored region. It seems that the Riishojgaard modification is acting, as anticipated, to honor the background moisture inversion structure, thereby limiting the ability of the analysis to produce excessive cloud thickness by artificially raising the cloud top.

Note that our earlier experiments with this marine stratocumulus case were a failure. In those experiments, we were using only

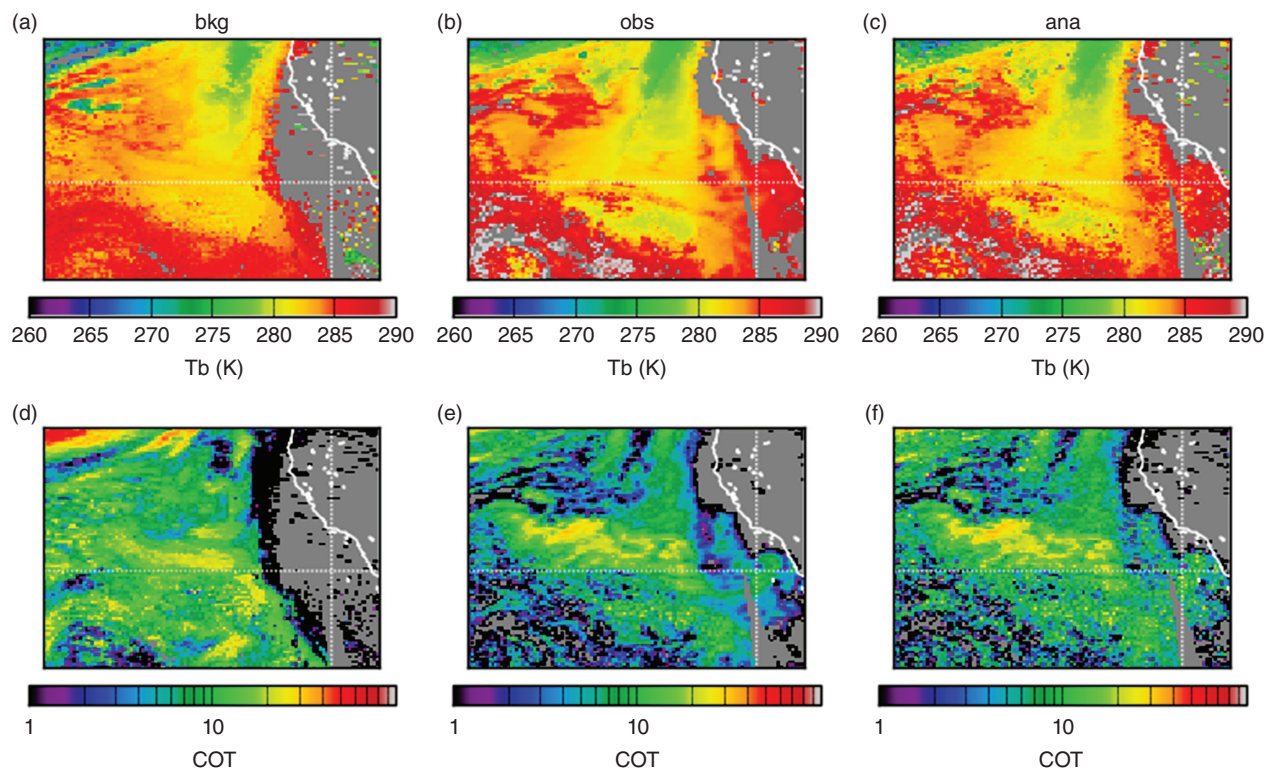


Figure 7. A case study of marine stratocumulus off the coast of California and Baja California (1 July 2011, at 2100 UTC). The upper panels (a–c) show brightness temperature T_b and the lower panels (d–f) all-sky COT (τ). The background (a,d) is absent the stratocumulus near the Southern or Baja Californian coast that is present in the observations (b,e). The analysis (c,f) is able to restore this stratocumulus, albeit with a slightly higher COT than observed. The somewhat noisier appearance of the analysis comes from the fact that each gridcolumn analysis is currently independent of its neighbours.

τ and p_c as observables, not T_b . The failure occurred because the MODIS-retrieved CO_2 -slicing cloud-top pressure is unreliable for low clouds (below 550 hPa) and was producing cloud-top pressures in the range 650–750 hPa, significantly above the marine inversion, which is below 900 hPa. The anatomy of the failure was as follows: the assimilation system was trying to produce a cloud with a top around 700 hPa (based on the erroneous MODIS p_c), but the background relative humidity at this height, well above the marine inversion, was only about 10%. Clearly, the prior in \bar{S} prevented the system from shifting the mean too far from 0.1, so the system instead produced a very positively skewed moisture PDF, with the upper tail just crossing saturation. The resulting small cloud fraction (less than 10%) near 700 hPa was able to raise the simulated cloud-top pressure towards the incorrect MODIS value, but at the expense of changing the observed cloud fraction of close to 100% to less than 10%! This illustrates the potential problems that can arise when assimilating erroneous data or alternatively using a forward model that does not simulate the actual observable well. In practice, had we made the cloud fraction an additional observable with a prior constraint, then the outcome would have been different. In that case, the highly skewed, low cloud fraction solution would have been impossible and the analysis would have remained near the background, which would have been preferable. As it was, we chose a better solution: we dropped the p_c observable in favour of T_b below 550 hPa.

6. Comparison between MODIS and SEVIRI assimilation

We ran an assimilation with a short 15 min window centred on 13:52:30 GMT of 1 July 2011, to compare the assimilation of MODIS *Aqua* cloud retrievals (as above) with assimilation of cloud retrievals made by the NASA Langley Cloud and Radiation Group from the SEVIRI instrument aboard the *Meteosat-9* geostationary platform (see Minnis *et al.*, 2008). (The above time is near the *Aqua* MODIS viewing of 0°N , 0°E , the subsatellite point of *Meteosat-9*.) The *Meteosat-9* SEVIRI retrievals have the advantage of large spatial coverage (including the whole of Europe, the Atlantic Ocean and all of Africa) at high temporal resolution

(every 15 min), compared with the twice-a-day overpasses of the *Aqua* and *Terra* polar-orbiting platforms. The nadir field of view is also reasonably small at 3 km, though not as small as the MODIS 1 km optical retrievals. From an assimilation point of view, the main difference is that only cloud optical thickness τ and brightness temperature T_b are assimilated from the SEVIRI retrievals, not CO_2 -slicing cloud-top pressure p_c as for MODIS. Thus, while the MODIS assimilation switches between T_b and p_c for cloudy pixels, depending on the value of p_c as described earlier and in Part 1, the SEVIRI assimilation always uses T_b when cloud is present.

Comparing Figures 8 and 9, it seems that assimilation of the two retrieval datasets is comparable (though the SEVIRI dataset has the obvious advantage of being available as a full-disc image every 15 min). An examination of the corresponding O–A PDFs (not shown) confirms this comparability, although the ability of the assimilation algorithm to remove unobserved cloud from the background appears to be a little better for MODIS than for SEVIRI assimilation. (Both experiments use the b7K64hC32 control with the tropopause (pTrop) and Riishojgaard (RiiS1) modifications discussed above.) Figure 10 summarizes the cross-validation of the MODIS *Aqua* and SEVIRI CDA results for the gridcolumn COT shown in Figures 8 and 9. It is evident that the two data sets are much closer to one another than the model background state is to either: mean $O_{\text{SEVIRI}} - O_{\text{MODIS}} \approx -1.9$, cf. mean $O_{\text{MODIS}} - B \approx -6.9$ and mean $O_{\text{SEVIRI}} - B \approx -8.8$. Furthermore, we see that assimilating SEVIRI cloud data strongly improves the validation against independent MODIS cloud data and vice versa. This provides further support for our Monte Carlo Bayesian CDA algorithm.

7. Validation using OMI

The above results provide substantial evidence that the Bayesian cloud data assimilation system is at least doing what it was designed to do: to drive the subcolumn model to a new state more consistent with the observations. We also have evidence, from the stratocumulus case study, that the CDA system is also able to

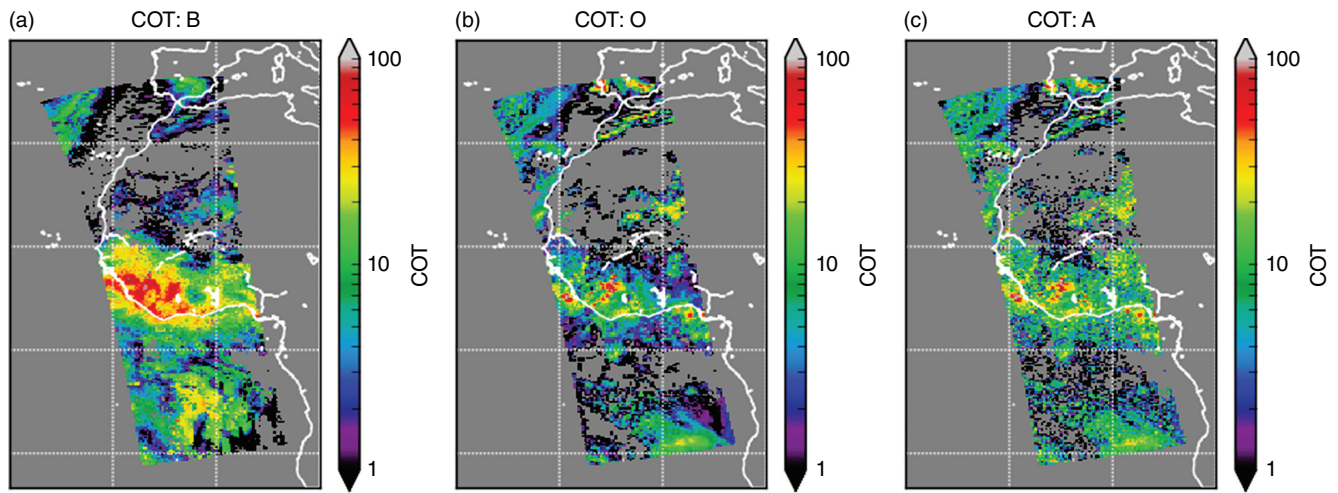


Figure 8. Cloud optical thickness for an assimilation of MODIS *Aqua* cloud data for a 15 min window centred on 13:52:30 GMT of 1 July 2011: (a) triangularized background, (b) observations, (c) CDA analysis. These results should be compared with the SEVIRI assimilation results in Figure 9.

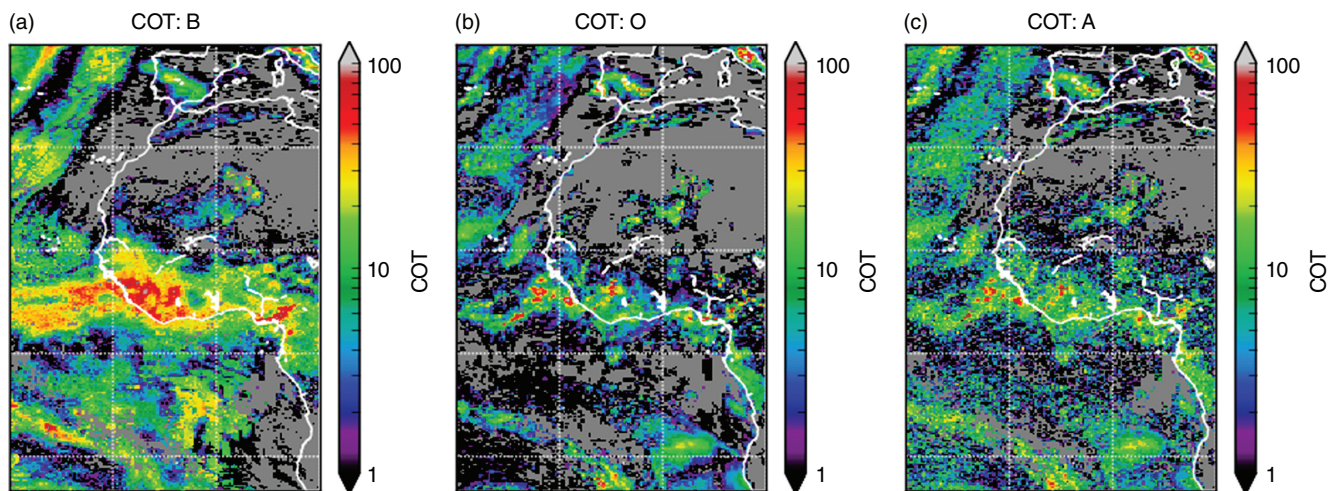


Figure 9. As in Figure 8, but for assimilation of cloud retrievals made by the NASA Langley Cloud and Radiation Group for *Meteosat-9* SEVIRI data (see text.)

handle particularly tricky cases where there is no linear sensitivity of the observations to the model. However, we also want a broader global validation of the improvement of the analyzed state compared with some independent, non-assimilated dataset. For this, we turn to the so-called Optical Centroid Pressure (OCP) retrieved from a rotational–Raman scattering (RRS) algorithm (see Joiner *et al.*, 2012, hereafter J12) applied to UV spectral measurements made by OMI aboard the polar-orbiting *Aura* satellite (at the rear of the so-called ‘A-Train’, 15 min behind *Aqua*).

The OCP is a measure of how far into a cloud one can see in the UV, which is much further than at IR wavelengths. Therefore, the OCP is much deeper inside a cloud than an IR window or CO₂-slicing-derived cloud-top pressure (CTP), which is typically very near the cloud top for all but very thin clouds. As such, the OCP is sensitive to both the CTP and the COT. That being said, we must realistically acknowledge that the Bayesian CDA system described in this series (Part 1) does not put a strong constraint on the vertical structure of the cloud, but mainly on the cloud top and the overall optical thickness. This limitation is particularly true for multi-layer clouds. Thus we cannot necessarily expect to achieve a major improvement in OCP simulations from the analyzed cloud state, but some improvement is hoped for.

For the observations, we use the OMI OMCLDRR product, which contains both the effective cloud fraction f_{eff} (‘CloudFractionforO3’) and optical centroid pressure P_{OCP} (‘CloudPressureforO3’) scientific datasets (SDS). The nadir field of view (FOV) is approximately 12 km along track and 24 km across track. We only include FOVs satisfying the following conditions, based on

guidance from J12 and the OMCLDRR README file provided with the data.

- (i) The OCP is valid if bits 0, 1, 2, 3, 4, 6, 7, 13, 14 and 15 of the ‘ProcessingQualityFlagsforO3’ SDS are all zero and if the P_{OCP} value itself has some positive value.
- (ii) f_{eff} is valid if P_{OCP} is and, further, we require $f_{\text{eff}} \geq 0.3$, since the algorithm is very noisy for small effective cloud fractions.
- (iii) We also exclude FOVs flagged as snow/ice by bit 5 of the above processing quality flag (pqf), since for these surface types f_{eff} is set to one and so P_{OCP} is approximate.
- (iv) We also exclude FOVs with radiance or irradiance errors according to bits 9 and 11 of the pqf.
- (v) Finally, we exclude FOVs possibly compromised by sea-glint under low cloud fraction conditions (namely $f_{\text{eff}} < 0.3$ and bit 4 of the ‘GroundPixelQualityFlags’ SDS set).

All FOVs satisfying these conditions are binned into the reduced longitude grid of the analysis and the mean (in-cloud) P_{OCP} is calculated for each accepted gridcolumn. In addition, only those gridcolumns for which the corresponding observed gridcolumn mean all-sky COT from *Aqua* is at least 5 are included in the analysis, since the retrieved OCP is believed to be less reliable for optically thinner clouds (see J12).

We study the day of 13 February 2007 and use a modified CDA control run that only assimilates *Aqua* (not *Terra*) observations, since OMI is on board *Aura*, which closely follows behind *Aqua*. A P_{OCP} is simulated for the triangularized background (B) and analyzed (A) states using the fast ‘R³S’ simulator of J12. Namely, P_{OCP} is evaluated according to Eq. (5) of

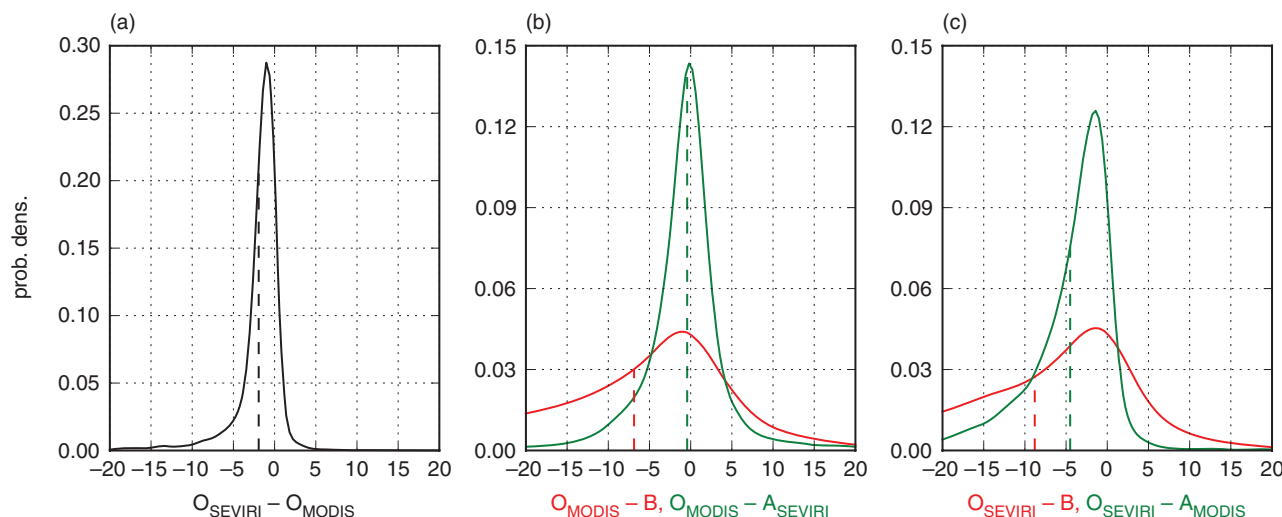


Figure 10. For the MODIS *Aqua* and SEVIRI gridcolumn COT of Figures 8 and 9, the following are shown. (a) The probability density of SEVIRI minus MODIS *Aqua* COT observations for gridcolumns sampled as cloudy by both instruments. (b) A cross-validation of the SEVIRI CDA (A_{SEVIRI}) against MODIS *Aqua* observations (O_{MODIS}), showing the probability density of CDA bias (green) compared with the density of background bias (red). Both densities use only gridcolumns sampled by *both* instruments and, for *each* density, both sides of the difference must be cloudy. (c) Conversely, a cross-validation of the MODIS *Aqua* CDA against SEVIRI observations. It is evident that MODIS *Aqua* and SEVIRI COT are much closer to one another than the model background state is to either. Furthermore, assimilating SEVIRI cloud data strongly improves the validation against MODIS and vice versa.

J12: $P_{\text{OCP}} = \sum_l \rho_l P_l / \sum_l \rho_l$, where P_l is the mid-point model layer pressure and ρ_l is the contribution to the cumulative UV cloud reflectance from layer l , as described in J12 (using the adding–doubling method and a conservative delta–Eddington two-stream calculation of layer reflectances and transmittances, with asymmetry factor $g = 0.8$).

Figure 11 shows P_{OCP} bias PDFs for O–B and O–A, globally and for different regions. The biases are in the in-cloud mean P_{OCP} for each gridcolumn in the specified region and for each of the eight assimilation time windows in the 24 h period.

Global mean biases are reduced by about 25% and the standard deviation in the biases by about 21%. Biases also appear to be more Gaussian after assimilation, which is helpful, given that these analyzed states will eventually become the background for a more traditional 3D- or 4D-variational data assimilation.

It is somewhat disappointing that the biases were not more strongly reduced, but, as stated earlier, the CDA method does not constrain cloud vertical structure strongly and also the fast R³S OCP simulator only provides an approximation to the retrieved OCP. To investigate the potential role of cloud vertical structure, we also studied the global mean bias and standard deviation reductions as a function of COT, τ . For moderately thin clouds ($5 \leq \tau \leq 10$), the global mean bias was reduced by 33% and the standard deviation by 24%. For clouds with $10 \leq \tau \leq 25$, these reductions are 27 and 23% and, for $\tau \geq 25$, only about 12 and 9%. Thus, the reductions do seem to deteriorate with increasing cloud thickness, especially for thick clouds ($\tau \geq 25$), which is consistent with the idea that cloud vertical structure will be more important for thick clouds. That being said, the addition of the ‘pTrop’ modification produced no significant change in the OCP bias reductions and the Riishojgaard modification produced a small degradation, so it is clear that the issue of constraint of vertical cloud structure is not yet well understood and needs further study. It should also be noted that the Riishojgaard modification still uses the same number of MTM trials M per chain point as the control, calculated based on the standard pressure scale-length method of Part 1 section 2.7.2. This should really be modified when using the Riishojgaard correlation modification, but it is not immediately clear to us how to do that.

Another potential issue affecting our OCP estimates could be our simplistic treatment of the cloud phase split and of mixed-phase clouds, as detailed in section 2.8 of Part 1. Certainly, the tropopause correction above would not improve this issue, since high clouds are mostly all ice. Nor would the Riishojgaard modification necessarily improve this matter, since many of the

buoyancy-inhibiting inversions capping the planetary boundary layer occur below the freezing level. In the final analysis, the treatment of phase is certainly a weak point that deserves further attention in subsequent studies.

8. Discussion and conclusions

Part 1 of this series presented a new Monte Carlo Bayesian method for constraining a complex statistical model of GCM sub-gridcolumn moisture variability using high-resolution MODIS cloud data. The method has strong application possibilities in cloud data assimilation and cloud parametrization testing and development, since it can be used to evaluate and improve background (prior) estimates of the state of a model gridcolumn and its subgrid parametrization parameters. We have chosen one particular gridcolumn statistical model for this study, namely skewed triangle PDFs of intra-layer moisture and a Gaussian copula to couple them in the vertical, but many other more or less complicated models can be conceived and intercompared using this method, for fidelity in simulating observed cloud data.

Part 2 has performed some basic testing of the new method, verifying that it does indeed reduce mean and standard deviation biases very significantly with respect to assimilated MODIS cloud optical depth, brightness temperature and cloud-top pressure. Of particular interest, a case study of marine stratocumulus off the Californian coast has demonstrated that the Monte Carlo method performs well in the especially difficult case where the background state is clear but cloudy observations exist. In traditional linearized data assimilation methods, a subsaturated background cannot produce clouds via any infinitesimal equilibrium perturbation, but the Monte Carlo approach allows non-gradient-based jumps into regions of non-zero cloud probability. In the example provided, the method was able to restore marine stratocumulus near the Californian coast, where the background state had a large clear swath.

One important application of the method is assimilation of cloud data into large-scale numerical weather prediction models. By showing that the method also easily assimilates geostationary SEVIRI retrievals, which are available every 15 min with huge spatial coverage, we have opened up the possibility of future significant assimilation of such retrievals to constrain the global moisture field, of which clouds are a significant marker.

The ultimate goal of this work is to produce a fully cycling data assimilation system, one in which the model total water PDF parameters are re-initialized with the values coming from

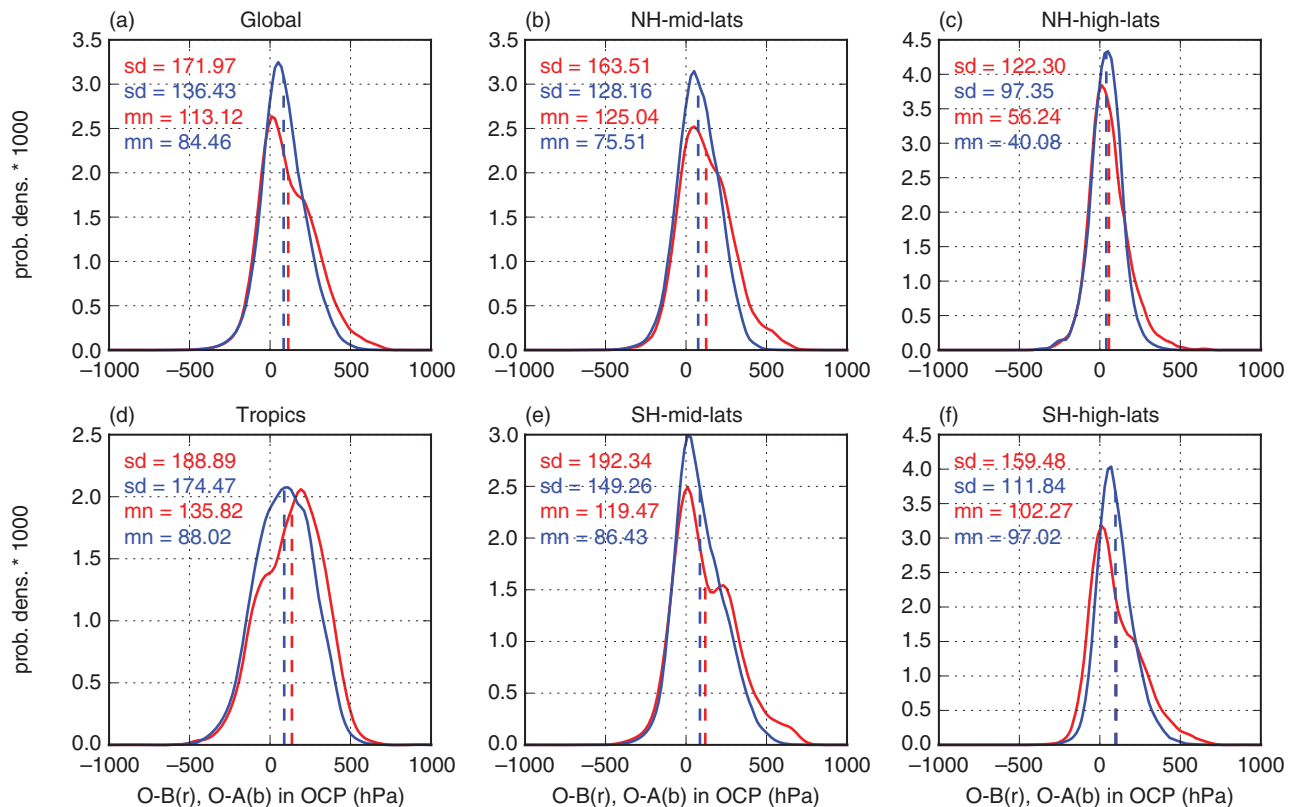


Figure 11. PDFs of the biases O–B (red) and O–A (blue) of in-cloud mean gridcolumn cloud optical centroid pressure P_{OCP} for a modified control run that assimilates only *Aqua* (not *Terra*) observations. Biases are shown for six regions (a)–(f), as noted, where the latitude boundaries between Tropics, mid and high latitudes are at 22.5 and 50°.

the ‘cloud analysis’, with the GCM producing a first guess for the next cloud analysis. To achieve this goal, we have structured the project with two distinct milestones: (i) development and assessment of the ‘cloud analysis’ step by means of MCMC and (ii) update of the PDF scheme in GEOS-5 to provide the time evolution of the triangular PDF parameters and to use the triangular PDF and subcolumn generation consistently in the radiation parametrization and throughout GEOS-5. This article, however, focuses on milestone (i). Moreover, the improved moisture/cloud state afforded by the MCMC algorithm is intended to provide a better background for the hybrid ensemble/variational algorithms in GEOS-5, as well as to assist in the development of proper observation operators for cloudy radiances (taking into consideration cloud overlapping and subgrid variability).

For computational feasibility, we have stayed away from a multivariate cloud analysis involving wind–mass coupling in the MCMC algorithm. Having obtained mass analysis increments, the correspond wind increments could, in principle, be generated by using the balance operators in our hybrid grid-point statistical interpolation (GSI) system or relegated altogether to the full meteorological analysis. The tacit goal of our approach is to extract cloud information from the wealth of visible and IR sensors as an intermediate step and to use this information to constrain better the assimilation of IR and microwave cloudy radiances in the main meteorological data assimilation system. What we present here is an incremental step in this direction.

Also, note that in many situations the maintenance of the analyzed cloud state requires support by other portions of the atmospheric state (surface fluxes, PBL mixing, subsidence, etc.). Many of these features are not constrained directly by the meteorological assimilation. Wind increments generated from mass increments by balance operators are often adiabatic, with minimal impact on subsidence (a largely unobserved aspect of the circulation). Therefore, we are not very optimistic that we will be able to resolve data retention issues with data assimilation only. Ultimately, model deficiencies will need to be addressed in order for cloud information to be retained. However, a cloud analysis

and its concurrent estimate of cloud radiative forcing can provide a very useful complementary dataset for atmospheric reanalyses, even when these data are only partially retained by the model.

One remaining difficulty for our method (and for other CDA methods as well) is the limited information content on cloud vertical structure in the chosen cloud observables, cloud-top pressure (or brightness temperature) and column optical depth. However, even these variables produce a significant improvement in the simulated rotational–Raman scattering cloud OCP against independent (non-assimilated) retrievals from the OMI instrument. Because one can see deeper into a cloud in the UV, the OCP is sensitive not only to the cloud-top pressure but also to cloud vertical structure. Clearly, the fidelity of the model’s background vertical structure, including boundary-layer inversions, will be important for this problem and it is important to honour such features in the assimilation process. It was found that a simple flow-dependent correlation function due to Riishojgaard (1998) provides some help in this respect, since it permits stronger decorrelations across moisture (or other) inversions in the background state. To improve the analyzed vertical structure of the cloud field further, future work could consider the addition of new observables, with stronger vertical sensitivity, to the Bayesian problem: OCP is one such observable, but other multispectral passive observations or active sensors can be considered as well. While current LIDAR and RADAR observations (e.g. *CALIPSO*, *CloudSat* and forthcoming sensors) contain the much needed vertical information, their coverage is usually too sparse to provide any lasting impact on global data assimilation systems.

Acknowledgements

This work was supported by NASA grants from the Modeling, Analysis, and Prediction program (solicitation NNNH08ZDA001N-MAP, proposal title: ‘Assimilation of *A-Train* satellite data for constraining a new PDF-based cloud parametrization in GEOS-5’, PI: Arlindo da Silva, and solicitation

NNH12ZDA001N-MAP, proposal title: ‘Using Near-Real Time Satellite Retrieved Cloud and Surface Properties to Validate and Improve GEOS-5 Analyses and Forecasts’, PIs: P. Minnis and M. Rienecker).

The authors thank Steven Platnick and Gala Wind for providing the MODIS cloud data and much useful insight into its properties and use, Pat Minnis and Rabindra Palikonda for the SEVIRI retrieval files used in the SEVIRI comparison and for guidance on their use and Joanna Joiner for providing the rotational–Raman scattering OCP fast-simulator code and guidance on the use of the OMI OMCLDRR dataset. The authors also thank Dr Chris Snyder and an anonymous reviewer for their reviews, which helped to improve this article substantially. Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center.

References

- Bacmeister JT, Suarez MJ, Robertson FR. 2006. Rain re-evaporation, boundary-layer/convection interactions and Pacific rainfall patterns in an AGCM. *J. Atmos. Sci.* **63**: 3383–3403.
- Joiner J, Vasilkov AP, Gupta P, Bhartia PK, Veeckind P, Sneep M, de Haan J, Polonsky I, Spurr R. 2012. Fast simulators for satellite cloud optical centroid pressure retrievals; evaluation of OMI cloud retrievals. *Atmos. Meas. Tech.* **5**: 529–545, doi: 10.5194/amt-5-529-2012.
- Larson VE. 2004. Prognostic equations for cloud fraction and liquid water, and their relation to filtered density functions. *J. Atmos. Sci.* **61**: 338–351.
- Larson VE, Wood R, Field PR, Golaz JC, Vonder Haar TH, Cotton WR. 2001. Small-scale and mesoscale variability of scalars in cloudy boundary layers: One-dimensional probability density functions. *J. Atmos. Sci.* **58**: 1978–1994.
- Minnis P, Nguyen L, Palikonda R, Heck PW, Spangenberg DA, Doelling DR, Ayers JK, Smith WL Jr, Khaiyer MM, Trepte QZ, Avey LA, Chang FL, Yost CR, Chee TL, Sun-Mack S. 2008. ‘Near-real time cloud retrievals from operational and research meteorological satellites’. In *Proceedings of SPIE Europe Remote Sensing*, 7107–2, 15–18 September, 2008. Cardiff, UK.
- Moorthi S, Suarez MJ. 1992. Relaxed Arakawa–Schubert, a parametrization of moist convection for general-circulation models. *Mon. Weather Rev.* **120**: 978–1002.
- Norris PM, da Silva AM. 2007. Assimilation of satellite cloud data into the GMAO finite volume data assimilation system using a parameter estimation method. *J. Atmos. Sci.* **64**: 3880–3895.
- Norris PM, da Silva AM. 2016. Monte Carlo Bayesian inference on a statistical model of sub-gridcolumn moisture variability using high-resolution cloud observations. Part 1: Method. *Q. J. R. Meteorol. Soc.* **142**: xxx–xxx (Part 1, this issue).
- Norris PM, Oreopoulos L, Hou AY, Tao WK, Zeng X. 2008. Representation of 3D heterogeneous cloud fields using copulas: Theory for water clouds. *Q. J. R. Meteorol. Soc.* **134**: 1843–1864, doi: 10.1002/qj.321.
- Ricard JL, Royer JF. 1993. A statistical cloud scheme for use in an AGCM. *Ann. Geophys.* **11**: 1095–1115.
- Rienecker MM, Suarez MJ, Gelaro R, Todling R, Bacmeister J, Liu E, Bosilovich MG, Schubert SD, Takacs L, Kim G-K, Bloom S, Chen J, Collins D, Conaty A, da Silva A, Gu W, Joiner J, Koster RD, Lucchesi R, Molod A, Owens T, Pawson S, Pegion P, Redder CR, Reichle R, Robertson FR, Ruddick AG, Sienkiewicz M, Woollen J. 2011. MERRA: NASA’s modern-era retrospective analysis for research and applications. *J. Clim.* **24**: 3624–3648, doi: 10.1175/JCLI-D-11-00015.1.
- Riishojgaard LP. 1998. A direct way of specifying flow-dependent background error correlations for meteorological analysis systems. *Tellus* **50A**: 42–57.
- Slingo JM. 1987. The development and verification of a cloud prediction scheme for the ECMWF model. *Q. J. R. Meteorol. Soc.* **113**: 899–927.
- Smith RNB. 1990. A scheme for predicting layer clouds and their water content in a general circulation model. *Q. J. R. Meteorol. Soc.* **116**: 435–460.
- Sommeria G, Deardorff J. 1977. Subgrid-scale condensation in models of nonprecipitating clouds. *J. Atmos. Sci.* **34**: 344–355.
- Tokioka T, Yamazaki K, Kitoh A, Ose T. 1988. The equatorial 30–60 day oscillation and the Arakawa–Schubert penetrative cumulus parametrization. *J. Meteorol. Soc. Jpn.* **66**: 883–901.
- Tompkins AM. 2002. A prognostic parametrization for the subgrid-scale variability of water vapour and clouds in large-scale models and its use to diagnose cloud cover. *J. Atmos. Sci.* **59**: 1917–1942.
- Xu KM, Randall DA. 1996. Evaluation of statistically based cloudiness parametrizations used in climate models. *J. Atmos. Sci.* **53**: 3103–3119.